arXiv:1610.03768v1 [math.GT] 12 Oct 2016

The high-dimensional cohomology of the moduli space of curves with level structures

Neil Fullarton* and Andrew Putman†

October 10, 2016

Abstract

We prove that the moduli space of curves with level structures has an enormous amount of rational cohomology in its cohomological dimension. As an application, we prove that the coherent cohomological dimension of the moduli space of curves is at least $g-2$. Well known conjectures of Looijenga would imply that this is sharp.

1 Introduction

Let Σ_g be a closed oriented genus g surface. The *mapping class group* of Σ_g, denoted Mod_g, is the group of isotopy classes of orientation-preserving diffeomorphisms of Σ_g. The group Mod_g lies at the crossroads of many areas of mathematics. One fundamental reason for this is that Mod_g is the orbifold fundamental group of the moduli space $\mathcal{M}_g$ of genus g Riemann surfaces. In fact, even more is true: as an orbifold, $\mathcal{M}_g$ is an Eilenberg–MacLane space for Mod_g, which implies in particular that

$$\mathrm{H}^*(\mathrm{Mod}_g;\mathbb{Q}) = \mathrm{H}^*(\mathcal{M}_g;\mathbb{Q}).$$

See [FarMar] for a survey of Mod_g and $\mathcal{M}_g$.

Stable cohomology. Let $\kappa_i \in \mathrm{H}^{2i}(\mathrm{Mod}_g;\mathbb{Q})$ be the i^{th} Miller–Mumford–Morita class. We then have a graded ring homomorphism $\mathbb{Q}[\kappa_1,\kappa_2,\ldots] \to \mathrm{H}^*(\mathrm{Mod}_g;\mathbb{Q})$, and the Mumford conjecture (proved by Madsen–Weiss [MadW]) says that this graded ring homomorphism is an isomorphism in degrees less than or equal to $\frac{2}{3}(g-1)$. Aside from some low-genus computations, no nontrivial elements of $\mathrm{H}^*(\mathrm{Mod}_g;\mathbb{Q})$ have been found outside this stable range. However, Harer–Zagier [HareZ] proved that the Euler characteristic of Mod_g is enormous, so there must exist vast amounts of unstable rational cohomology.

Level structures. The cohomology of finite-index subgroups of Mod_g (or, equivalently, finite covers of $\mathcal{M}_g$) is also of interest. For $\ell \geqslant 2$, the *level ℓ congruence subgroup* of Mod_g, denoted $\mathrm{Mod}_g(\ell)$, is the kernel of the action of Mod_g on $\mathrm{H}_1(\Sigma_g;\mathbb{Z}/\ell)$. It fits into a short exact sequence

$$1 \longrightarrow \mathrm{Mod}_g(\ell) \longrightarrow \mathrm{Mod}_g \longrightarrow \mathrm{Sp}_{2g}(\mathbb{Z}/\ell) \longrightarrow 1.$$

The symplectic group appears here because the action of Mod_g on $\mathrm{H}_1(\Sigma_g;\mathbb{Z}/\ell)$ preserves the algebraic intersection pairing. The associated finite cover of $\mathcal{M}_g$ is the moduli space $\mathcal{M}_g(\ell)$ of genus g curves equipped with a full level ℓ structure (i.e. a basis for the ℓ-torsion in their Jacobian). A conjecture of the second author (see [Pu, §1] for a discussion) asserts that

$$\mathrm{H}_k(\mathrm{Mod}_g(\ell);\mathbb{Q}) \cong \mathrm{H}_k(\mathrm{Mod}_g;\mathbb{Q}) \qquad (g \gg k).$$

This holds for $k=1$ by work of Hain [Hai] and for $k=2$ by work of the second author [Pu].

*Supported in part by an AMS-Simons grant and an MSRI postdoctoral fellowship
†Supported in part by NSF grant DMS-1255350 and the Alfred P. Sloan Foundation

Cohomological dimension. The main topic of this paper is what happens outside the stable range. Harer [Hare2] proved that the virtual cohomological dimension (vcd) of Mod_g is $4g-5$, so $\mathrm{H}^i(\mathrm{Mod}_g;\mathbb{Q}) = 0$ for $i > 4g-5$. The first place where one might hope to find some unstable cohomology is thus in degree $4g-5$. However, a theorem proved independently by Morita–Sakasai–Suzuki [MorSaSu] and by Church–Farb–Putman [ChuFarPu] says that $\mathrm{H}^{4g-5}(\mathrm{Mod}_g;\mathbb{Q}) = 0$. This might seem to contradict the fact that the vcd of Mod_g is $4g-5$. However, the definition of the vcd of a group makes use not only of ordinary cohomology, but also cohomology with respect to arbitrary twisted coefficient systems. Harer's theorem thus only asserts that there exists some $\mathbb{Q}[\mathrm{Mod}_g]$-module M (necessarily nontrivial, in light of [MorSaSu, ChuFarPu]) such that $\mathrm{H}^{4g-5}(\mathrm{Mod}_g;M) \neq 0$.

Main theorem. This brings us to our main theorem, which says that in contrast to what conjecturally happens in the stable range, the group $\mathrm{Mod}_g(\ell)$ has an enormous amount of cohomology in its vcd.

Theorem A. *Fix $g, \ell \geqslant 2$ and let p be a prime dividing ℓ. Then*

$$\dim_{\mathbb{Q}} \mathrm{H}^{4g-5}(\mathrm{Mod}_g(\ell);\mathbb{Q}) \geqslant \frac{|\mathrm{Sp}_{2g}(\mathbb{F}_p)|}{g(p^{2g}-1)} = \frac{1}{g}p^{2g-1}\prod_{k=1}^{g-1}(p^{2k}-1)p^{2k-1}.$$

Remark 1.1. Our lower bound is super-exponential; its leading term is $\frac{1}{g}p^{\binom{2g}{2}}$. To give an idea of how quickly it grows, the following are some special cases:

$$\begin{array}{ll} \dim_{\mathbb{Q}} \mathrm{H}^3(\mathrm{Mod}_2(2);\mathbb{Q}) \geqslant 24 & \dim_{\mathbb{Q}} \mathrm{H}^3(\mathrm{Mod}_2(3);\mathbb{Q}) \geqslant 216 \\ \dim_{\mathbb{Q}} \mathrm{H}^7(\mathrm{Mod}_3(2);\mathbb{Q}) \geqslant 11520 & \dim_{\mathbb{Q}} \mathrm{H}^7(\mathrm{Mod}_3(3);\mathbb{Q}) \geqslant 4199040 \\ \dim_{\mathbb{Q}} \mathrm{H}^{11}(\mathrm{Mod}_4(2);\mathbb{Q}) \geqslant 92897280 & \dim_{\mathbb{Q}} \mathrm{H}^{11}(\mathrm{Mod}_4(3);\mathbb{Q}) \geqslant 6685442749440. \end{array}$$

Remark 1.2. Our lower bound is almost certainly not sharp. One difficulty with improving it is that as described below, we exploit a connection to the Tits building for the group $\mathrm{SL}_n(\mathbb{F}_p)$. Presumably better results could be obtained by studying the obvious analogue of this building for the finite group $\mathrm{SL}_n(\mathbb{Z}/\ell)$, but for ℓ not prime this is poorly understood.

Remark 1.3. In his 1986 paper [Hare2] (see p. 175), Harer asserts that "it is possible to show" that $\mathrm{H}^{4g-5}(\mathrm{Mod}_g(\ell);\mathbb{Q}) \neq 0$. However, he never published a proof of this.

Application to algebraic geometry. Theorem A has an interesting application to the algebraic geometry of $\mathcal{M}_g$, which recall is the (coarse) moduli space of genus g Riemann surfaces. We begin with the following conjecture of Looijenga [FabLo].

Conjecture 1.4 (Looijenga)**.** *For $g \geqslant 2$, the quasiprojective variety $\mathcal{M}_g$ can be covered by $(g-1)$ open affine subsets.*

For example, this conjecture asserts that $\mathcal{M}_2$ is itself affine, which is a consequence of the fact that every genus 2 Riemann surface is hyperelliptic. More generally, Fontanari–Pascolutti [FoPa] proved that Conjecture 1.4 holds for $2 \leqslant g \leqslant 5$.

Conjecture 1.4 would imply a bound on the coherent cohomological dimension of $\mathcal{M}_g$, which is defined as follows. If X is a variety, then the *coherent cohomological dimension* of X, denoted $\mathrm{CohCD}(X)$, is the maximum value of k such that there exists some coherent sheaf $\mathcal{F}$ on X with $\mathrm{H}^k(X;\mathcal{F}) \neq 0$. The coherent cohomological dimension of a variety reflects interesting geometric properties of the variety. For example, Serre [Se] proved that $\mathrm{CohCD}(X) = 0$ if and only if X is an affine variety. See [Hart] for more information on coherent cohomological dimension.

Since $\mathcal{M}_g$ is separated, the intersection of two affine open subsets of $\mathcal{M}_g$ is itself an affine open subset. Thus if Conjecture 1.4 were true and $\mathcal{F}$ were a coherent sheaf on $\mathcal{M}_g$, then we could apply the Mayer–Vietoris spectral sequence to the cover of $\mathcal{M}_g$ given by Conjecture 1.4 to deduce that $\mathrm{H}^k(\mathcal{M}_g;\mathcal{F}) = 0$ for $k > g-2$. In other words, Conjecture 1.4 would imply that $\mathrm{CohCD}(\mathcal{M}) \leqslant g-2$.

Using our main theorem (Theorem A), we will prove the following, which asserts that this conjectural upper bound is sharp.

Theorem B. *For $g \geqslant 2$, we have $\mathrm{CohCD}(\mathcal{M}_g) \geqslant g-2$ with equality for $2 \leqslant g \leqslant 5$.*

Remark 1.5. This implies that $\mathcal{M}_g$ cannot be covered with fewer than $(g-1)$ open affine subsets. This was already known. Indeed, it follows from work of Chaudhuri [Cha], who gave a lower bound on the cohomological excess of $\mathcal{M}_g$ (which is defined using constructible sheaves).

Remark 1.6. The only paper we are aware of concerning upper bounds for things related to $\mathrm{CohCD}(\mathcal{M}_g)$ is recent work of Mondello [Mon], who proved that the Dolbeault cohomological dimension of $\mathcal{M}_g$ is at most $2g-2$. The Dolbeault cohomological dimension of a complex analytic variety X is the maximal k such that there exists a holomorphic vector bundle $\mathcal{B}$ on X such that $\mathrm{H}^k(\mathcal{M}_g;\mathcal{B}) \neq 0$. As will be clear from our argument below, our work also establishes a lower bound of $g-2$ on the Dolbeault cohomological dimension of $\mathcal{M}_g$.

Proof of Theorem B. The derivation of Theorem B from Theorem A is so simple that we give it here. We wish to thank Eduard Looijenga for explaining this argument to us. Fix some $\ell \geqslant 3$, so $\mathcal{M}_g(\ell)$ is smooth. The projection $\mathcal{M}_g(\ell) \to \mathcal{M}_g$ is a finite surjective map, so $\mathrm{CohCD}(\mathcal{M}_g(\ell)) = \mathrm{CohCD}(\mathcal{M}_g)$ (see [Hart, Proposition 1.1]). It is thus enough to prove that $\mathrm{CohCD}(\mathcal{M}_g(\ell)) \geqslant g-2$. Assume for the sake of contradiction that this is false. The Hodge–de Rham spectral sequence for $\mathcal{M}_g(\ell)$ converges to $\mathrm{H}^*(\mathcal{M}_g(\ell);\mathbb{C})$ and has

$$E_1^{pq} = \mathrm{H}^p(\mathcal{M}_g(\ell);\Omega^q).$$

Since the complex dimension of $\mathcal{M}_g(\ell)$ is $3g-3$, we have $\Omega^q = 0$ for $q \geqslant 3g-2$, and thus

$$E_1^{pq} = 0 \qquad (q \geqslant 3g-2). \tag{1}$$

Moreover, since Ω^q is a coherent sheaf on $\mathcal{M}_g(\ell)$ our assumption that $\mathrm{CohCD}(\mathcal{M}_g(\ell)) < g-2$ implies that

$$E_1^{pq} = \mathrm{H}^p(\mathcal{M}_g(\ell);\Omega^q) = 0 \qquad (p \geqslant g-2). \tag{2}$$

From (1) and (2), we deduce that $E_1^{pq} = 0$ whenever $p+q = 4g-5$. This implies that $\mathrm{H}^{4g-5}(\mathcal{M}_g(\ell);\mathbb{C}) = 0$, contradicting Theorem A. The fact that we have equality for $2 \leqslant g \leqslant 5$ follows from the aforementioned theorem of Fontanari–Pascolutti [FoPa] asserting that Conjecture 1.4 holds for $2 \leqslant g \leqslant 5$.

Proof outline for Theorem A. Our proof of Theorem A has four steps.

1. First, we use the fact that the mapping class group satisfies Bieri–Eckmann duality [Hare2] to translate the theorem into an assertion about the action of $\mathrm{Mod}_g(\ell)$ on the Steinberg module for the mapping class group, i.e. the unique nonzero homology group of the curve complex.

2. Next, we study this action by constructing a novel surjective homomorphism from the Steinberg module for the mapping class group to a vector space $\mathrm{St}^{\mathrm{ns}}_{2g}(\mathbb{F}_p)$ that is a quotient of the Steinberg module $\mathrm{St}_{2g}(\mathbb{F}_p)$ for the finite group $\mathrm{SL}_{2g}(\mathbb{F}_p)$, i.e. the unique nonzero homology group of this finite group's Tits building.

It follows from the previous two steps that the dimension of $\mathrm{H}^{4g-5}(\mathrm{Mod}_g(\ell);\mathbb{Q})$ is at least $\dim_{\mathbb{Q}} \mathrm{St}^{\mathrm{ns}}_{2g}(\mathbb{F}_p)$. At this point, one might think that we have at least proved that $\mathrm{H}^{4g-5}(\mathrm{Mod}_g(\ell);\mathbb{Q}) \neq 0$. However, there is a problem – from its definition, it is not clear that $\mathrm{St}^{\mathrm{ns}}_{2g}(\mathbb{F}_p) \neq 0$. The next two steps analyze this vector space.

3. The third step is representation-theoretic. As a prelude to analyzing $\mathrm{St}^{\mathrm{ns}}_{2g}(\mathbb{F}_p)$, we show how to decompose the restriction of the $\mathrm{SL}_{2g}(\mathbb{F}_p)$-representation $\mathrm{St}_{2g}(\mathbb{F}_p)$ to the subgroup $\mathrm{Sp}_{2g}(\mathbb{F}_p)$, i.e. we construct a branching rule between these two different classical groups (see Remark 4.2 below).

4. Finally, to use the third step to show that $\mathrm{St}^{\mathrm{ns}}_{2g}(\mathbb{F}_p)$ is nonzero and satisfies the bound in Theorem A, we apply classical theorems concerning the partition function and exponential generating functions (some of which go back to Euler).

Remark 1.7. One might expect that the Steinberg module for the finite group $\mathrm{Sp}_{2g}(\mathbb{F}_p)$ would appear here rather than the Steinberg module for $\mathrm{SL}_{2g}(\mathbb{F}_p)$. We tried to do this initially, but were unsuccessful. Indeed, every map from the Steinberg module for the mapping class group to the Steinberg module for $\mathrm{Sp}_{2g}(\mathbb{F}_p)$ that we were able to concoct ended up being the zero map. What is more, our lower bound on the dimension of $\mathrm{H}^{4g-5}(\mathrm{Mod}_g(\ell);\mathbb{Q})$ is significantly larger than the dimension of the Steinberg module for $\mathrm{Sp}_{2g}(\mathbb{F}_p)$, namely p^{g^2}, so it seems unlikely that one could use the Steinberg module for $\mathrm{Sp}_{2g}(\mathbb{F}_p)$ to prove a theorem as strong as Theorem A.

Outline. Our proof of Theorem A is contained in §2 – 5. The first part of the proof is in §2, which as described above uses Bieri–Eckmann duality to connect Theorem A to the Steinberg module for the mapping class group. In §3, we describe how to connect the Steinberg module for the mapping class group to the Steinberg module for the special linear group. Next, in §4 we construct the branching rule discussed above. Finally, in §5 we establish the lower bound in Theorem A.

Acknowledgments. We wish to thank Eduard Looijenga for asking us whether Theorem A held and explaining why it implied Theorem B. We also want to thank Thomas Church, Benson Farb, Martin Kassabov, Dan Margalit, John Wiltshire-Gordon, and the MathOverflow user "Lucia" for helpful comments.

2 Bieri–Eckmann duality for the mapping class group

In this section, we translate Theorem A into a statement (Theorem A′ below) concerning the Steinberg module for the mapping class group.

Ignoring orbifold issues, the space $\mathcal{M}_g$ is a smooth orientable manifold. If it were compact, then it would satisfy Poincaré duality. Unfortunately, it is not compact. It therefore only satisfies Poincaré–Lefshetz duality with a contribution coming from the "boundary" (i.e. at infinity). This is best described in terms of the group Mod_g. Harer [Hare2] proved

that Mod_g satisfies a version of Poincaré–Lefshetz duality called virtual Bieri–Eckmann duality. Recalling that the vcd of Mod_g is $4g-5$, this duality take the form

$$\mathrm{H}^{4g-5-i}(\mathrm{Mod}_g;\mathbb{Q}) = \mathrm{H}_i(\mathrm{Mod}_g;\mathrm{St}(\Sigma_g)) \qquad (i \geqslant 0). \tag{3}$$

Here the dualizing module $\mathrm{St}(\Sigma_g)$ is the Steinberg module for Mod_g, which is defined as follows. The *curve complex* for Σ_g, denoted $\mathcal{C}_g$, is the simplicial complex whose k-simplices are collections $\{\gamma_0,\dots,\gamma_k\}$ of distinct isotopy classes of non-nullhomotopic simple closed curves on Σ_g that can be realized disjointly. One way of thinking of $\mathcal{C}_g$ is that it is a combinatorial model for all the ways that a genus g Riemann surfaces can degenerate to a stable nodal surface. The group Mod_g acts on $\mathcal{C}_g$, and Harer [Hare2] proved that $\mathcal{C}_g$ is homotopy equivalent to an infinite wedge of $(2g-2)$-dimensional spheres. The *Steinberg module* for Mod_g is defined to be

$$\mathrm{St}(\Sigma_g) \coloneqq \widetilde{\mathrm{H}}_{2g-2}(\mathcal{C}_g;\mathbb{Q}).$$

In the duality (3), this measures the contribution "at infinity" to cohomology.

Since $\mathrm{Mod}_g(\ell)$ is a finite-index subgroup of Mod_g, it also satisfies Bieri–Eckmann duality with the same dualizing module $\mathrm{St}(\Sigma_g)$. We thus have

$$\mathrm{H}^{4g-5-i}(\mathrm{Mod}_g(\ell);\mathbb{Q}) = \mathrm{H}_i(\mathrm{Mod}_g(\ell);\mathrm{St}(\Sigma_g)) \qquad (i \geqslant 0).$$

Recall that if G is a group and M is a G-module, then $\mathrm{H}_0(G;M) = M_G$. Here M_G denotes the *coinvariants* of G acting on M, i.e. the quotient of M by the submodule spanned by $\{m - g(m) \mid m \in M,\, g \in G\}$. To prove Theorem A, we must prove that

$$\mathrm{H}_0(\mathrm{Mod}_g(\ell);\mathrm{St}(\Sigma_g)) = (\mathrm{St}(\Sigma_g))_{\mathrm{Mod}_g(\ell)}$$

is large, i.e. we must construct $\mathrm{Mod}_g(\ell)$-invariant homomorphisms from $\mathrm{St}(\Sigma_g)$ whose targets are large. More precisely, Theorem A is equivalent to the following theorem.

Theorem A′. *Fix $g,\ell \geqslant 2$, and let p be a prime dividing ℓ. Then there exists a $\mathrm{Mod}_g(\ell)$-invariant surjective homomorphism $\psi\colon \mathrm{St}(\Sigma_g) \to V$ such that V is a vector space over $\mathbb{Q}$ satisfying*

$$\dim_{\mathbb{Q}} V = \frac{|\,\mathrm{Sp}_{2g}(\mathbb{F}_p)|}{g(p^{2g}-1)}.$$

3 Mapping to the nonseparated building

In this section, we construct the homomorphism ψ whose existence is asserted by Theorem A′. Later sections will prove that its target is as large as claimed by that theorem. Though we will not use any of his results explicitly, our point of view in this section is heavily influenced by Broaddus's paper [Broa]. To avoid degenerate situations, we will fix some $g \geqslant 2$.

Arcs. To construct ψ, we need to understand generators for $\mathrm{St}(\Sigma_g) = \widetilde{\mathrm{H}}_{2g-2}(\mathcal{C}_g;\mathbb{Q})$. For this, we will use an alternate model for $\mathcal{C}_g$ that was constructed by Harer. Fix a basepoint $* \in \Sigma_g$. The *arc complex* on Σ_g, denoted $\mathcal{AC}_g^1$, is the simplicial complex whose k-simplices are collections $\{\alpha_0,\dots,\alpha_k\}$ of distinct homotopy classes of unoriented $*$-based simple closed curves on Σ_g that can be realized to be disjoint away from $*$. The superscript 1 in $\mathcal{AC}_g^1$

indicates that there is a single basepoint. Harer [Hare1, Theorem 1.5] proved that $\mathcal{AC}_g^1$ is contractible. See [Hat] for a beautiful short proof of this. Next, define $\mathcal{AC}_g^1(\infty)$ to be the subcomplex of $\mathcal{AC}_g^1$ consisting of simplices $\{\alpha_0, \ldots, \alpha_k\}$ such that some component of Σ_g cut along $\alpha_0 \cup \cdots \cup \alpha_k$ is non-contractible. Harer [Hare1, Lemma 1.7] proved that $\mathcal{AC}_g^1(\infty)$ is homotopy equivalent to $\mathcal{C}_g$.

Remark 3.1. What Harer actually proves in the above reference is that $\mathcal{AC}_g^1(\infty)$ is homotopy equivalent to the curve complex on a once-punctured surface; however, another theorem of Harer from [Hare2] is that the curve complex on a once-punctured surface is homotopy equivalent to the curve complex on a closed surface. See [KLeiSc] and [HatV] for alternate proofs of this.

Lemma 3.2. *For $g \geqslant 2$, we have* $\mathrm{St}(\Sigma_g) \cong \mathrm{H}_{2g-1}(\mathcal{AC}_g^1, \mathcal{AC}_g^1(\infty); \mathbb{Q})$.

Proof. Since $\mathcal{AC}_g^1$ is contractible, the long exact sequence in relative homology implies that $\mathrm{H}_i(\mathcal{AC}_g^1(\infty); \mathbb{Q}) \cong \mathrm{H}_{i+1}(\mathcal{AC}_g^1, \mathcal{AC}_g^1(\infty); \mathbb{Q})$. The lemma now follows from the fact that $\mathcal{AC}_g^1(\infty)$ is homotopy equivalent to $\mathcal{C}_g$. □

Remark 3.3. The description of $\mathrm{St}(\Sigma_g)$ in Lemma 3.2 slightly obscures the Mod_g-action, which goes as follows (cf. Remark 3.1). Consider a mapping class which is represented by an orientation-preserving diffeomorphism $f\colon \Sigma_g \to \Sigma_g$. Isotoping f, we can assume that $f(*) = *$, and thus f induces an automorphism of $\mathcal{AC}_g^1$ that preserves $\mathcal{AC}_g^1(\infty)$. Under the isomorphism in Lemma 3.2, the induced action of f on $\mathrm{H}_{2g-1}(\mathcal{AC}_g^1, \mathcal{AC}_g^1(\infty); \mathbb{Q})$ is precisely the action of the mapping class associated to f on $\mathrm{St}(\Sigma_g)$. In particular, this action on homology is independent of all of our choices.

Lemma 3.4. *For $g \geqslant 1$, the $(2g-2)$-skeletons of $\mathcal{AC}_g^1$ and $\mathcal{AC}_g^1(\infty)$ are equal.*

Proof. Consider a k-simplex $\sigma = \{\alpha_0, \ldots, \alpha_k\}$ of $\mathcal{AC}_g^1$ that does not lie in $\mathcal{AC}_g^1(\infty)$. Fix an orientation on each α_i. Since all the components of the complement of the α_i are simply-connected, the homology classes of the α_i must span $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$. This implies that $k+1 \geqslant 2g$, as desired. □

Lemma 3.5. *For $g \geqslant 2$, we have* $\mathrm{St}(\Sigma_g) \cong \mathrm{coker}(C_{2g}(\mathcal{AC}_g^1; \mathbb{Q}) \xrightarrow{\partial} C_{2g-1}(\mathcal{AC}_g^1, \mathcal{AC}_g^1(\infty); \mathbb{Q}))$.

Proof. Immediate from Lemmas 3.2 and 3.4. □

The building. Our next goal is to construct the target of ψ. This target should be $\mathrm{Mod}_g(\ell)$-invariant, and the obvious $\mathrm{Mod}_g(\ell)$-invariant object associated to Mod_g is the homology group $\mathrm{H}_1(\Sigma_g; \mathbb{Z}/\ell)$. Letting p be a prime dividing ℓ, the quotient $\mathrm{H}_1(\Sigma_g; \mathbb{F}_p) \cong \mathbb{F}_p^{2g}$ of $\mathrm{H}_1(\Sigma_g; \mathbb{Z}/\ell)$ is also invariant under $\mathrm{Mod}_g(\ell)$. To explain how we will use this to construct a target for ψ, we need to introduce Tits buildings. Fix some $n \geqslant 2$. The *Tits building* associated to $\mathrm{SL}_n(\mathbb{F}_p)$, denoted $\mathcal{T}_n(\mathbb{F}_p)$, is the simplicial complex whose r-simplices are flags

$$0 \subsetneq V_0 \subsetneq V_1 \subsetneq \cdots \subsetneq V_r \subsetneq \mathbb{F}_p^n.$$

The simplicial complex $\mathcal{T}_n(\mathbb{F}_p)$ is $(n-2)$-dimensional, and the Solomon–Tits theorem ([So]; see also [Brow, Theorem IV.5.2]) says that in fact $\mathcal{T}_n(\mathbb{F}_p)$ is homotopy equivalent to a wedge of $(n-2)$-dimensional spheres. The *Steinberg module* for $\mathrm{SL}_n(\mathbb{F}_p)$ is defined to be

$$\mathrm{St}_n(\mathbb{F}_p) \coloneqq \widetilde{\mathrm{H}}_{n-2}(\mathcal{T}_n(\mathbb{F}_p); \mathbb{Q}).$$

This is one of the most important representations of $\mathrm{SL}_n(\mathbb{F}_p)$; for instance, it is the unique nontrivial irreducible representation of $\mathrm{SL}_n(\mathbb{F}_p)$ whose dimension is a power of p.

Apartments. The Solomon–Tits theorem also gives generators for $\mathrm{St}_n(\mathbb{F}_p)$. Let $\mathfrak{X}_n$ be the simplicial complex whose r-simplices are increasing sequences

$$\varnothing \subsetneq I_0 \subsetneq I_1 \subsetneq \cdots \subsetneq I_r \subsetneq \{1, \ldots, n\}$$

of sets. The simplicial complex $\mathfrak{X}_n$ is isomorphic to the boundary of the barycentric subdivision of an $(n-1)$-simplex, and thus is homeomorphic to an oriented $(n-2)$-sphere; let $[\mathfrak{X}_n]$ be its fundamental class. Associated to an ordered sequence $B = (\vec{v}_1, \ldots, \vec{v}_n)$ of n nonzero vectors in $\mathbb{F}_p^n$, there is a simplicial map $\mathfrak{X}_n \to \mathcal{T}_n$ taking the vertex $\varnothing \subsetneq I \subsetneq \{1, \ldots, n\}$ to $\langle \vec{v}_i \mid i \in I \rangle$. The image under this map of $[\mathfrak{X}_n] \in \widetilde{\mathrm{H}}(\mathfrak{X}_n; \mathbb{Q})$ is the *apartment class* $\mathfrak{A}_B \in \mathrm{St}_n(\mathbb{F}_p) = \widetilde{\mathrm{H}}_{n-2}(\mathcal{T}_n(\mathbb{F}_p); \mathbb{Q})$. The Solomon–Tits theorem asserts that $\mathrm{St}_n(\mathbb{F}_p)$ is generated by apartment classes.

Properties of apartments. The following result of Lee–Szczarba [LeeSz] summarizes a number of basic properties of apartment classes.

Theorem 3.6 ([LeeSz, Theorem 3.1])**.** *Fix $n, p \geqslant 2$ such that p is prime. The following then hold.*

1. *For an apartment class $\mathfrak{A}_B$ in $\mathrm{St}_n(\mathbb{F}_p)$, we have $\mathfrak{A}_B \neq 0$ if and only if B is a basis for $\mathbb{F}_p^n$.*
2. *For an apartment class $\mathfrak{A}_B$ in $\mathrm{St}_n(\mathbb{F}_p)$ and for a permutation σ of $\{1, \ldots, n\}$, we have $\mathfrak{A}_B = (-1)^{|\sigma|}\mathfrak{A}_{\sigma(B)}$. Here $\sigma(B)$ is defined in the obvious way.*
3. *For an apartment class $\mathfrak{A}_B$ in $\mathrm{St}_n(\mathbb{F}_p)$ and for nonzero scalars $c_1, \ldots, c_n \in \mathbb{F}_p^*$, write $B = (\vec{v}_1, \ldots, \vec{v}_n)$ and let $B' = (c_1\vec{v}_1, \ldots, c_n\vec{v}_n)$. Then $\mathfrak{A}_B = \mathfrak{A}_{B'}$.*
4. *For an ordered sequence $C = (\vec{w}_1, \ldots, \vec{w}_{n+1})$ of nonzero vectors in $\mathbb{F}_p^n$, define $C_i = (\vec{w}_1, \ldots, \widehat{\vec{w}_i}, \ldots, \vec{w}_{n+1})$. Then $\sum_{i=1}^{n+1}(-1)^i\mathfrak{A}_{C_i} = 0$.*

What is more, these generate all relations between apartment classes in $\mathrm{St}_n(\mathbb{F}_p)$ in the sense that $\mathrm{St}_n(\mathbb{F}_p)$ is the $\mathbb{Q}$-vector space with generators the set of formal symbols

$$\{\mathfrak{A}_B \mid B \text{ an ordered sequence of } n \text{ nonzero vectors in } \mathbb{F}_p^n\}$$

and relations the four relations listed above.

Remark 3.7. Lee–Szczarba stated their result differently than we have, but the two formulations are equivalent. The only non-obvious relation among those above is the fourth one; see Figure 1 for an explanation of it.

The map on chains. If V is a vector space over $\mathbb{Q}$, then Lemma 3.5 implies that a homomorphism

$$\psi\colon \mathrm{St}(\Sigma_g) \to V$$

can be constructed by writing down a homomorphism

$$\phi\colon \mathrm{C}_{2g-1}(\mathcal{AC}_g^1, \mathcal{AC}_g^1(\infty); \mathbb{Q}) \to V$$

that vanishes on the image of

$$\partial\colon \mathrm{C}_{2g}(\mathcal{AC}_g^1; \mathbb{Q}) \to \mathrm{C}_{2g-1}(\mathcal{AC}_g^1, \mathcal{AC}_g^1(\infty); \mathbb{Q}).$$

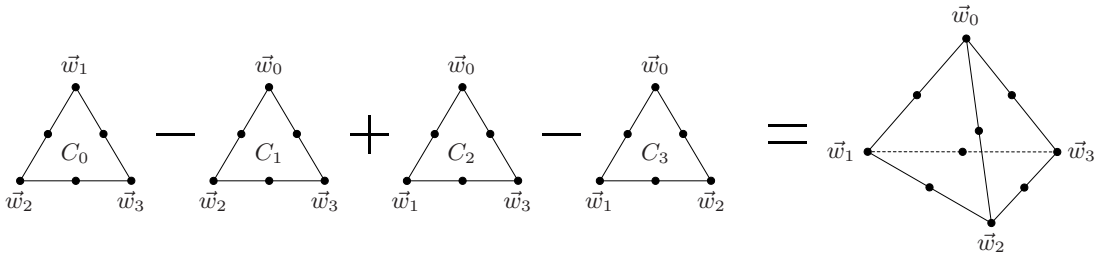

Figure 1: *As we illustrate here in the case $n = 3$, the apartment classes $\mathfrak{A}_{C_i}$ can be placed on the boundary of an n-dimensional simplex. In the picture, the vertices labeled with the vectors $\vec{w}_i$ represent the lines spanned by the $\vec{w}_i$ while the unlabeled vertices represent the 2-dimensional subspaces spanned by the vectors on their two neighbors.*

Recall that p is a fixed prime dividing ℓ. Fix an identification of $\mathrm{H}_1(\Sigma_g; \mathbb{F}_p)$ with $\mathbb{F}_p^{2g}$ that takes the algebraic intersection form on $\mathrm{H}_1(\Sigma_g; \mathbb{F}_p)$ to the standard symplectic form on $\mathbb{F}_p^{2g}$. Given an oriented closed curve γ on Σ_g, let $[\gamma]_p$ be the associated element of $\mathrm{H}_1(\Sigma_g; \mathbb{F}_p) = \mathbb{F}_p^{2g}$. We can define a map $\widetilde{\phi}\colon \mathrm{C}_{2g-1}(\mathcal{AC}_g^1; \mathbb{Q}) \to \mathrm{St}_{2g}(\mathbb{F}_p)$ as follows.

- Let σ be an oriented $(2g-1)$-simplex of $\mathcal{AC}_g^1$. Write $\sigma = \{\alpha_0, \ldots, \alpha_{2g-1}\}$, and orient each curve α_i in an arbitrary way. If $[\alpha_i]_p = 0$ for some i, then define $\widetilde{\phi}(\sigma) = 0$. Otherwise, set $B = ([\alpha_0]_p, \ldots, [\alpha_{2g-1}]_p)$ and define $\widetilde{\phi}(\sigma) = \mathfrak{A}_B$.

The second and third relations in Theorem 3.6 show that $\widetilde{\phi}$ does not depend on any of the choices we have made. Moreover, the first relation in Theorem 3.6 shows that $\widetilde{\phi}(\sigma) = 0$ if there is any non-simply connected component in the complement of the α_i, so $\widetilde{\phi}$ vanishes on $\mathrm{C}_{2g-1}(\mathcal{AC}_g^1(\infty))$ and thus descends to a map

$$\phi\colon \mathrm{C}_{2g-1}(\mathcal{AC}_g^1, \mathcal{AC}_g^1(\infty); \mathbb{Q}) \to \mathrm{St}_{2g}(\mathbb{F}_p).$$

A problem. Unfortunately, ϕ does not vanish on the image of ∂. Consider an oriented $2g$-simplex $\eta = \{\beta_0, \ldots, \beta_{2g}\}$ of $\mathcal{AC}_g^1$. There are two cases where $\phi(\partial(\eta)) = 0$.

- If all the $[\beta_i]_p$ are nonzero, then ϕ takes $\partial(\eta)$ to a relation in $\mathrm{St}_{2g}(\mathbb{F}_p)$ like in the fourth relation in Theorem 3.6, and in particular $\phi(\partial(\eta)) = 0$.
- If more than one of the $[\beta_i]_p$ are zero, then the first relation in Theorem 3.6 implies ϕ takes every simplex in $\partial(\eta)$ to zero, so $\phi(\partial(\eta)) = 0$.

However, if exactly one of the $[\beta_i]_p$ is zero, then $\phi(\partial(\eta)) \neq 0$. To understand precisely what is going on, assume that $[\beta_j]_p = 0$ for some j but that $[\beta_i]_p \neq 0$ for all $i \neq j$. It follows that β_j is a separating simple closed curve (if it were nonseparating, then it would represent a primitive element of $\mathrm{H}_1(\Sigma_g; \mathbb{Z})$, and hence its homology class would remain nonzero when reduced modulo p). Letting S_1 and S_2 be the subsurfaces on either side of β_j, we obtain a splitting

$$\mathrm{H}_1(\Sigma_g; \mathbb{F}_p) = \mathrm{H}_1(S_1; \mathbb{F}_p) \oplus \mathrm{H}_1(S_2; \mathbb{F}_p)$$

that is orthogonal with respect to the algebraic intersection form. Moreover, for each $i \neq j$, the homology class $[\beta_i]_p$ lies in either $\mathrm{H}_1(S_1; \mathbb{F}_p)$ or $\mathrm{H}_1(S_2; \mathbb{F}_p)$.

Separated apartments. This motivates the following definition. An apartment $\mathfrak{A}_B$ of $\mathrm{St}_{2g}(\mathbb{F}_p)$ is a *separated apartment* if there exists a nontrivial splitting $\mathbb{F}_p^{2g} = S_1 \oplus S_2$ that is orthogonal with respect to the standard symplectic form on $\mathbb{F}_p^{2g}$ such that each $\vec{v}$ occurring

in B lies in either S_1 or S_2. We remark that the orthogonality of S_1 and S_2 implies that the symplectic form on $\mathbb{F}_p^{2g}$ restricts to a symplectic form on each S_i. The *separated subspace* of $\mathrm{St}_{2g}(\mathbb{F}_p)$, denoted $\mathrm{St}_{2g}^{\mathrm{sep}}(\mathbb{F}_p)$, is the subspace spanned by separated apartments. While $\mathrm{St}_{2g}^{\mathrm{sep}}(\mathbb{F}_p)$ is not preserved by the action of $\mathrm{SL}_{2g}(\mathbb{F}_p)$, it is preserved by the subgroup $\mathrm{Sp}_{2g}(\mathbb{F}_p)$ of $\mathrm{SL}_{2g}(\mathbb{F}_p)$. By the previous paragraph, for each $2g$-simplex η of $\mathcal{AC}_g^1$, either $\phi(\partial(\eta)) = 0$ or $\phi(\partial(\eta)) \in \mathrm{St}_{2g}^{\mathrm{sep}}(\mathbb{F}_p)$. Define

$$\mathrm{St}_{2g}^{\mathrm{ns}}(\mathbb{F}_p) = \mathrm{St}_{2g}(\mathbb{F}_p)/\,\mathrm{St}_{2g}^{\mathrm{sep}}(\mathbb{F}_p),$$

so $\mathrm{St}_{2g}^{\mathrm{ns}}(\mathbb{F}_p)$ is an $\mathrm{Sp}_{2g}(\mathbb{F}_p)$-representation. We deduce that ϕ descends to a well-defined map

$$\psi\colon\ \mathrm{St}(\Sigma_g) = \mathrm{coker}(\partial) \longrightarrow \mathrm{St}_{2g}^{\mathrm{ns}}(\mathbb{F}_p).$$

By its construction, ψ is equivariant under the action of the whole mapping class group Mod_g and is invariant under the action of the subgroup $\mathrm{Mod}_g(\ell)$.

Putting everything together. The main result of this section is the following, which is part of the assertion of Theorem A′.

Proposition 3.8. *Fix $g, \ell \geqslant 2$, and let p be a prime dividing ℓ. Then the map*

$$\psi\colon\ \mathrm{St}(\Sigma_g) \longrightarrow \mathrm{St}_{2g}^{ns}(\mathbb{F}_p)$$

constructed above is surjective.

Proof. For $x \in \mathrm{St}_{2g}(\mathbb{F}_p)$, write $[\![x]\!]$ for the image of x in $\mathrm{St}_{2g}^{\mathrm{ns}}(\mathbb{F}_p)$. Also, let $\omega(\cdot,\cdot)$ be the standard symplectic form on $\mathbb{F}_p^{2g}$. Our proof will have two steps.

Step 1. *The vector space $\mathrm{St}_{2g}^{ns}(\mathbb{F}_p)$ is spanned by the set of $[\![\mathfrak{A}_B]\!]$ such that $B = (\vec{v}_1, \dots, \vec{v}_{2g})$ satisfies the following two conditions:*

- *$\omega(\vec{v}_i, \vec{v}_{i+1}) = 1$ for $1 \leqslant i < 2g$, and*
- *$\omega(\vec{v}_i, \vec{v}_j) = 0$ for $1 \leqslant i, j \leqslant 2g$ such that $|i - j| > 1$.*

Proof of Step 1. The proof of this step is inspired by the proof of [Broa, Proposition 4.6]. For $1 \leqslant k \leqslant 2g$, let V_k be the subspace of $\mathrm{St}_{2g}^{\mathrm{ns}}(\mathbb{F}_p)$ spanned by the set of $[\![\mathfrak{A}_B]\!]$ such that $B = (\vec{v}_1, \dots, \vec{v}_{2g})$ satisfies the following three conditions:

- $\omega(\vec{v}_i, \vec{v}_{i+1}) = 1$ for $1 \leqslant i < k$, and
- $\omega(\vec{v}_i, \vec{v}_j) = 0$ for $1 \leqslant i, j \leqslant k$ such that $|i - j| > 1$, and
- $\omega(\vec{v}_i, \vec{v}_j) = 0$ for $1 \leqslant i < k$ and $k < j \leqslant 2g$.

We will prove that V_k spans $\mathrm{St}_{2g}^{\mathrm{ns}}(\mathbb{F}_p)$ by induction on k; the case $k = 2g$ is precisely the claim we are trying to prove. The base case $k = 1$ is trivial, so assume that $1 < k \leqslant 2g$ and that V_{k-1} spans $\mathrm{St}_{2g}^{\mathrm{ns}}(\mathbb{F}_p)$. Consider an apartment $\mathfrak{A}_B$ such that B satisfies the three conditions above for $k-1$, and thus $[\![\mathfrak{A}_B]\!] \in V_{k-1}$. It is enough to prove that $[\![\mathfrak{A}_B]\!] \in V_k$. Since the symplectic form $\omega(\cdot,\cdot)$ is nondegenerate, we can find $\vec{w} \in \mathbb{F}_p^{2g}$ such that the following two conditions hold:

- $\omega(\vec{w}, \vec{v}_1) = 1$, and

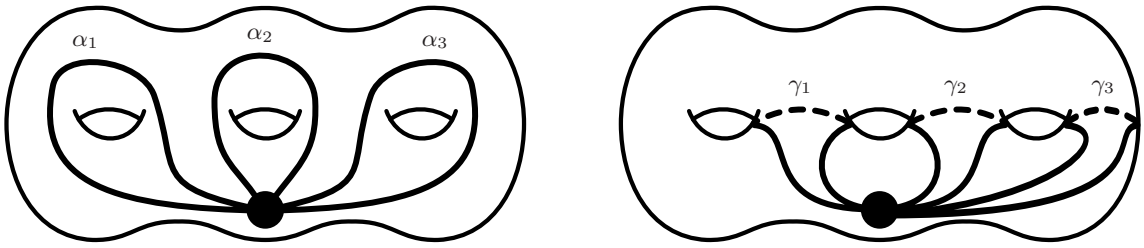

Figure 2: *A picture for $g = 3$ of the simplex $\sigma = \{\alpha_1, \gamma_1, \alpha_2, \gamma_2, \ldots, \alpha_g, \gamma_g\}$ of $\mathcal{AC}_g^1$ used in the proof of Proposition 3.8. The pattern for higher g is evident.*

- $\omega(\vec{w}, \vec{v}_i) = 0$ for $2 \leqslant i \leqslant 2g$.

Define $C = (\vec{w}, \vec{v}_1, \vec{v}_2, \ldots, \vec{v}_{2g})$, and for $1 \leqslant i \leqslant 2g+1$, let C_i be the result of deleting the i^{th} vector in C. The fourth relation in Theorem 3.6 says that

$$\sum_{i=1}^{2g+1} (-1)^i \mathfrak{A}_{C_i} = 0 \tag{4}$$

in $\mathrm{St}_{2g}(\mathbb{F}_p)$. We now make the following observations.

- $C_1 = B$.
- For $2 \leqslant i \leqslant k$, we have $\mathfrak{A}_{C_i} \in \mathrm{St}_{2g}^{\mathrm{sep}}(\mathbb{F}_p)$, so $[\![\mathfrak{A}_{C_i}]\!] = 0$. We remark that some of these apartment classes are actually 0.
- For $k < i \leqslant 2g+1$, we have $\mathfrak{A}_{C_i} \in V_k$

Combining these three observations with (4), we see that $[\![\mathfrak{A}_B]\!] \in V_k$, as desired. □

To prove Proposition 3.8, it is enough to show that the image of ψ contains all the generators identified in Step 1, which is the content of our next step.

Step 2. *Consider an apartment $\mathfrak{A}_B$ such that $B = (\vec{v}_1, \ldots, \vec{v}_{2g})$ satisfies the two conditions stated in Step 1. Then there exists an element $x \in \mathrm{St}(\Sigma_g)$ such that $\psi(x) = [\![\mathfrak{A}_B]\!]$.*

Proof of Step 2. The group $\mathrm{Sp}_{2g}(\mathbb{F}_p)$ acts transitively on ordered bases B for $\mathbb{F}_p^{2g}$ that satisfy the two conditions in Step 1. Since the map $\mathrm{Mod}_g \to \mathrm{Sp}_{2g}(\mathbb{F}_p)$ coming from the action on $\mathrm{H}_1(\Sigma_g; \mathbb{F}_p)$ is surjective and the homomorphism ψ intertwines the mapping class group action on $\mathrm{St}(\Sigma_g)$ and the symplectic group action on $\mathrm{St}_{2g}^{\mathrm{ns}}(\mathbb{F}_p)$, we see that it is enough to find an element $x \in \mathrm{St}(\Sigma_g)$ such that $\psi(x) = [\![\mathfrak{A}_C]\!]$, where $C = (\vec{w}_1, \ldots, \vec{w}_{2g})$ is *some* ordered basis for $\mathbb{F}_p^{2g}$ satisfying the two conditions in Step 1 (with $\vec{w}_i$ swapped for $\vec{v}_i$). Examining the construction of ψ, we see that we can take x to be the image in $\mathrm{St}(\Sigma_g)$ of the $(2g-1)$-simplex

$$\sigma = \{\alpha_1, \gamma_1, \alpha_2, \gamma_2, \ldots, \alpha_g, \gamma_g\}$$

of $\mathcal{AC}_g^1$ depicted in Figure 2. □

This completes the proof of Proposition 3.8. □

To derive Theorem A′ from Proposition 3.8, it is enough to prove the following proposition.

Proposition 3.9. *Consider $g, p \geqslant 1$ such that p is prime. Then*

$$\dim_{\mathbb{Q}} \mathrm{St}_{2g}^{ns}(\mathbb{F}_p) = \frac{|\mathrm{Sp}_{2g}(\mathbb{F}_p)|}{g(p^{2g}-1)}.$$

We will prove Proposition 3.9 in §5. This is preceded by a careful study of the structure of $\mathrm{St}_{2g}(\mathbb{F}_p)$ as an $\mathrm{Sp}_{2g}(\mathbb{F}_p)$ representation in §4.

4 Decomposing the nonseparated building

Fix a prime p. All vector spaces in this section are finite-dimensional vector spaces over $\mathbb{F}_p$. In preparation for proving Proposition 3.9 in §5, this section is devoted to understanding the $\mathrm{Sp}_{2g}(\mathbb{F}_p)$-module structure of $\mathrm{St}_{2g}(\mathbb{F}_p)$. Our main result is Theorem 4.1 below.

Abstract vector spaces. Before stating our main result, we need some preliminaries. Given a vector space V, let $\mathcal{T}(V)$ be the Tits building and $\mathrm{St}(V)$ be the Steinberg module associated to V. Thus if $n = \dim_{\mathbb{F}_p}(V)$, then $\mathrm{St}(V) = \widetilde{\mathrm{H}}_{n-2}(\mathcal{T}(V); \mathbb{Q})$. If V is a symplectic vector space, then let $\mathrm{St}^{\mathrm{sep}}(V)$ denote the separated subspace of $\mathrm{St}(V)$ as defined in §3. Also, define $\mathrm{St}^{\mathrm{ns}}(V) = \mathrm{St}(V)/\mathrm{St}^{\mathrm{sep}}(V)$. By Maschke's theorem, we can embed the $\mathrm{Sp}(V)$-representation $\mathrm{St}^{\mathrm{ns}}(V)$ into $\mathrm{St}(V)$ such that we have an isomorphism

$$\mathrm{St}(V) = \mathrm{St}^{\mathrm{sep}}(V) \oplus \mathrm{St}^{\mathrm{ns}}(V) \tag{5}$$

of $\mathrm{Sp}(V)$-representations.

Unordered tensor products. Given vector spaces $V_1, \ldots, V_k$, the *unordered tensor product* of the V_i, denoted $\mathrm{Sym}(V_1, \ldots, V_k)$, is defined as follows. Let $\widehat{\mathrm{Sym}}(V_1, \ldots, V_k)$ be the direct sum over all elements σ of the k-element symmetric group of $V_{\sigma(1)} \otimes \cdots \otimes V_{\sigma(k)}$. Then $\mathrm{Sym}(V_1, \ldots, V_k)$ is the coinvariants of $\widehat{\mathrm{Sym}}(V_1, \ldots, V_k)$ under the action of the k-element symmetric group. As its name suggests, $\mathrm{Sym}(V_1, \ldots, V_k)$ is exactly like the ordinary tensor product, but without a distinguished ordering of the factors; in particular, we have an equality

$$\mathrm{Sym}(V_1, \ldots, V_k) = \mathrm{Sym}(V_{\sigma(1)}, \ldots, V_{\sigma(k)})$$

for all elements σ of the k-element symmetric group. The composition

$$V_1 \otimes \cdots \otimes V_k \hookrightarrow \widehat{\mathrm{Sym}}(V_1, \ldots, V_k) \to \mathrm{Sym}(V_1, \ldots, V_k)$$

is an isomorphism. Given elements $\vec{v}_i \in V_i$ for $1 \leqslant i \leqslant k$, we will write $\mathrm{Sym}(\vec{v}_1, \ldots, \vec{v}_k)$ for the image in $\mathrm{Sym}(V_1, \ldots, V_k)$ of $\vec{v}_1 \otimes \cdots \otimes \vec{v}_k$.

Symplectic splittings. A *symplectic splitting* of $\mathbb{F}_p^{2g}$ is an unordered set $S = \{S_1, \ldots, S_k\}$ of nonzero subspaces of $\mathbb{F}_p^{2g}$ that are pairwise orthogonal to each other under the symplectic form such that $\mathbb{F}_p^{2g} = S_1 \oplus \cdots \oplus S_k$. We remark that this implies that the symplectic form on $\mathbb{F}_p^{2g}$ restricts to a symplectic form on each S_i. Given a symplectic splitting $S = \{S_1, \ldots, S_k\}$ of $\mathbb{F}_p^{2g}$, we define

$$\mathrm{St}(S) = \mathrm{Sym}(\mathrm{St}(S_1), \ldots, \mathrm{St}(S_k))$$

and

$$\mathrm{St}^{\mathrm{ns}}(S) = \mathrm{Sym}(\mathrm{St}^{\mathrm{ns}}(S_1), \ldots, \mathrm{St}^{\mathrm{ns}}(S_k)).$$

There is a projection $\mathrm{St}(S) \to \mathrm{St}^{\mathrm{ns}}(S)$ whose kernel is the span of the subspaces

$$\{\mathrm{Sym}(\mathrm{St}^{\mathrm{sep}}(S_i), \mathrm{St}(S_1), \ldots, \widehat{\mathrm{St}(S_i)}, \ldots, \mathrm{St}(S_k)) \mid 1 \leqslant i \leqslant k\}$$

of $\mathrm{St}(S)$. Denote this kernel by $\mathrm{St}^{\mathrm{sep}}(S)$. We can regard $\mathrm{St}^{\mathrm{ns}}(S)$ as a subspace of $\mathrm{St}(S)$ via the embedding used in (5), and with this convention we have

$$\mathrm{St}(S) = \mathrm{St}^{\mathrm{sep}}(S) \oplus \mathrm{St}^{\mathrm{ns}}(S). \tag{6}$$

Let $\mathrm{Sp}(S)$ be the stabilizer of S in $\mathrm{Sp}_{2g}(\mathbb{F}_p)$. The vector spaces $\mathrm{St}(S)$ and $\mathrm{St}^{\mathrm{sep}}(S)$ and $\mathrm{St}^{\mathrm{ns}}(S)$ are representations of $\mathrm{Sp}(S)$, and (6) is a decomposition as $\mathrm{Sp}(S)$-representations.

Main theorem, statement. We can now state our main theorem. Let $\mathfrak{S}_g$ be the set of all symplectic splittings of $\mathbb{F}_p^{2g}$ (including the 1-element splitting $\{\mathbb{F}_p^{2g}\}$).

Theorem 4.1. *Let $g \geqslant 1$ and $p \geqslant 2$ be such that p is prime. Then we have an isomorphism*

$$\mathrm{St}_{2g}(\mathbb{F}_p) \cong \bigoplus_{S \in \mathfrak{S}_g} \mathrm{St}^{ns}(S)$$

of $\mathrm{Sp}_{2g}(\mathbb{F}_p)$-representations.

Remark 4.2. Theorem 4.1 can be stated in more representation-theoretic terms: letting $U \subset \mathfrak{S}_g$ be a set containing a single representative of each $\mathrm{Sp}_{2g}(\mathbb{F}_p)$-orbit, we have

$$\bigoplus_{S \in \mathfrak{S}_g} \mathrm{St}^{\mathrm{ns}}(S) \cong \bigoplus_{S \in U} \mathrm{Ind}_{\mathrm{Sp}(S)}^{\mathrm{Sp}_{2g}(\mathbb{F}_p)} \mathrm{St}^{\mathrm{ns}}(S).$$

We do not know if these induced representations are irreducible or not.

Product maps. Before we prove Theorem 4.1, we will construct product maps $\mathrm{St}(S) \to \mathrm{St}_{2g}(\mathbb{F}_p)$ and projection maps $\mathrm{St}_{2g}(\mathbb{F}_p) \to \mathrm{St}(S)$. We begin with the product maps. If $S = \{S_1, \ldots, S_k\}$ is a symplectic splitting $\mathbb{F}_p^{2g}$ and $\mathfrak{A}_{B_i}$ is an apartment class in $\mathrm{St}(S_i)$ for each for each $1 \leqslant i \leqslant k$, then let $B_1 \cdots B_k$ denote the ordered sequence of nonzero vectors in $\mathbb{F}_p^{2g}$ obtained by concatenating the B_i, so $\mathfrak{A}_{B_1 \cdots B_k}$ is an apartment class in $\mathbb{F}_p^{2g}$. We then have the following lemma.

Lemma 4.3. *Let $g \geqslant 1$ and $p \geqslant 2$ be such that p is prime. Let $S = \{S_1, \ldots, S_k\}$ be a symplectic splitting of $\mathbb{F}_p^{2g}$. Then there is a linear map $\iota_S \colon \mathrm{St}(S) \to \mathrm{St}_{2g}(\mathbb{F}_p)$ such that if $\mathfrak{A}_{B_i}$ is an apartment class in $\mathrm{St}(S_i)$ for $1 \leqslant i \leqslant k$, then*

$$\iota_S(Sym(\mathfrak{A}_{B_1}, \ldots, \mathfrak{A}_{B_k})) = \mathfrak{A}_{B_1 \cdots B_k}.$$

Proof. Set $n_i = \dim_{\mathbb{F}_p}(S_i)$, so $n_1 + \cdots + n_k = 2g$. Since each S_i is a symplectic subspace of $\mathbb{F}_p^{2g}$, each n_i is even. Since permuting the vectors forming an apartment class only changes the apartment class by a sign (see Theorem 3.6), this implies that the indicated formula for ι_S does not depend on the ordering of the B_i, and is thus well-defined.

Theorem 3.6 gives presentations for the $\mathrm{St}(S_i)$ and hence for $\mathrm{St}(S)$. We must check that the indicated formula take the relations for $\mathrm{St}(S)$ to relations for $\mathrm{St}_{2g}(\mathbb{F}_p)$. The only nonobvious relation is the one coming from the fourth relation in Theorem 3.6, which goes as follows. To simplify our notation, we will check this relation in the $\mathrm{St}(S_1)$-factor of $\mathrm{St}(S)$;

the other verifications are similar. Let C be an ordered sequence of n_1+1 nonzero vectors in S_1, and for $2 \leqslant i \leqslant k$ let B_i be an ordered sequence of n_i nonzero vectors in S_i. For $1 \leqslant i \leqslant n_1+1$, let C_i be the result of deleting the i^{th} vector from C. We then get a relation

$$\sum_{i=1}^{n_1+1} (-1)^i \mathrm{Sym}(\mathfrak{A}_{C_i}, \mathfrak{A}_{B_2}, \ldots, \mathfrak{A}_{B_k}) = 0$$

in $\mathrm{St}(S)$. We must check that ι_S takes this to a relation in $\mathrm{St}_{2g}(\mathbb{F}_p)$, i.e. we must check that

$$\sum_{i=1}^{n_1+1} (-1)^i \mathfrak{A}_{C_i B_2 \cdots B_k} = 0. \tag{7}$$

Let D be the result of concatenating C with B_2 through B_k, so D is a sequence of $2g+1$ nonzero vectors in $\mathbb{F}_p^{2g}$. For $1 \leqslant i \leqslant 2g+1$, let D_i be the result of deleting the i^{th} vector from D. We thus have $D_i = C_i B_2 \cdots B_k$ for $1 \leqslant i \leqslant n_1+1$. The fourth relation in Theorem 3.6 says that

$$0 = \sum_{i=1}^{2g+1} (-1)^i \mathfrak{A}_{D_i} = \left(\sum_{i=1}^{n_1+1} (-1)^i \mathfrak{A}_{C_i B_2 \cdots B_k} \right) + \left(\sum_{i=n_1+2}^{2g+1} (-1)^i \mathfrak{A}_{D_i} \right). \tag{8}$$

For $n_1+2 \leqslant i \leqslant 2g+1$ the vectors in D_i do not span $\mathbb{F}_p^{2g}$, so the first relation in Theorem 3.6 implies that $\mathfrak{A}_{D_i} = 0$. Plugging this into (8), we obtain (7), as desired. □

Projection maps, ordered. We now turn to the projection maps discussed before the statement of Lemma 4.3. We will do this in two steps. The first step requires some definitions. An *ordered symplectic splitting* of $\mathbb{F}_p^{2g}$ is an ordered sequence $\widehat{S} = (S_1, \ldots, S_k)$ of distinct subspaces of $\mathbb{F}_p^{2g}$ such that $S \coloneqq \{S_1, \ldots, S_k\}$ is a symplectic splitting of $\mathbb{F}_p^{2g}$. We will say that $\widehat{S}$ is an *ordering* of S. Next, if $\widehat{S} = (S_1, \ldots, S_k)$ is an ordered symplectic splitting of $\mathbb{F}_p^{2g}$, then an ordered sequence B of $2g$ vectors in $\mathbb{F}_p^{2g}$ is *weakly compatible* with $\widehat{S}$ if we can write $B = B_1 \cdots B_k$ such that $B_1 \cdots B_i$ is a basis of $S_1 \oplus \cdots \oplus S_i$ for all $1 \leqslant i \leqslant k$. In that case, we will say that $B_1 \cdots B_k$ is the *factorization* of B associated to S. Also, the *projection* of B to $\widehat{S}$ is the sequence $(\overline{B}_1, \ldots, \overline{B}_k)$, where for $1 \leqslant i \leqslant k$ the sequence $\overline{B}_i$ is obtained by applying the projection

$$S_1 \oplus \cdots \oplus S_{i-1} \oplus S_i \longrightarrow S_i$$

to each vector in B_i.

Lemma 4.4. *Let $g \geqslant 1$ and $p \geqslant 2$ be such that p is prime. Let $\widehat{S} = (S_1, \ldots, S_k)$ be an ordered symplectic splitting of $\mathbb{F}_p^{2g}$. Then there is a linear map $\pi_{\widehat{S}} \colon \mathrm{St}_{2g}(\mathbb{F}_p) \to \mathrm{St}(S)$ such that the following properties hold. Let $\mathfrak{A}_B$ be an apartment class in $\mathrm{St}_{2g}(\mathbb{F}_p)$.*

- *If B is weakly compatible with $\widehat{S}$ and $(\overline{B}_1, \ldots, \overline{B}_k)$ is its projection to $\widehat{S}$, then*

$$\pi_{\widehat{S}}(\mathfrak{A}_B) = \mathit{Sym}(\mathfrak{A}_{\overline{B}_1}, \ldots, \mathfrak{A}_{\overline{B}_k}).$$

- *If no reordering of the vectors in B is weakly compatible with $\widehat{S}$, then $\pi_{\widehat{S}}(\mathfrak{A}_B) = 0$.*

Proof. For $1 \leqslant i \leqslant k$, let $n_i = \dim_{\mathbb{F}_p}(S_i)$, so $n_1 + \cdots + n_k = 2g$. Also, for $0 \leqslant i \leqslant k$ define $T_i = S_1 \oplus \cdots \oplus S_i$, so

$$0 = T_0 \subsetneq T_1 \subsetneq \cdots \subsetneq T_k = \mathbb{F}_p^{2g}.$$

Define X to be the full subcomplex of $\mathrm{St}_{2g}(\mathbb{F}_p)$ spanned by vertices V (i.e. nonzero proper subspaces of $\mathbb{F}_p^{2g}$) such that there exists some $0 \leqslant i < k$ and some nonzero proper subspace W of S_{i+1} such that $V = T_i \oplus W$. Topologically, X is the join $\mathcal{T}(S_1) * \cdots * \mathcal{T}(S_k)$, and thus is a $(2g-k-1)$-dimensional simplicial complex which is $(2g-k-2)$-connected and satisfies

$$\mathrm{H}_{2g-k-1}(X;\mathbb{Q}) \cong \mathrm{St}(S_1) \otimes \cdots \otimes \mathrm{St}(S_k) \cong \mathrm{Sym}(\mathrm{St}(S_1), \ldots, \mathrm{St}(S_k)).$$

Our goal is thus to construct a homomorphism

$$\pi_{\widehat{S}} \colon \mathrm{H}_{2g-2}(\mathcal{T}_{2g}(\mathbb{F}_p);\mathbb{Q}) \longrightarrow \mathrm{H}_{2g-k-1}(X;\mathbb{Q}). \tag{9}$$

Before we do this, we make an observation concerning X. Consider a q-simplex μ of X. By definition, μ is a flag

$$0 \subsetneq V_0 \subsetneq V_1 \subsetneq \cdots \subsetneq V_q \subsetneq \mathbb{F}_p^{2g}$$

such that each V_i is of the form $T_{j_i} \oplus W_i$ for some $0 \leqslant j_i \leqslant k-1$ and some nonzero $W_i \subsetneq S_{j_i+1}$. There is thus a unique way of inserting $T_1, \ldots, T_{k-1}$ into η in such a way that it remains a flag. The result is a $(q+k-1)$-simplex of $\mathcal{T}_{2g}(\mathbb{F}_p)$. This establishes a bijection between the q-simplices of X and the $(q+k-1)$-simplices of $\mathcal{T}_{2g}(\mathbb{F}_p)$ whose associated flags contain $T_1, \ldots, T_{k-1}$.

We now turn to constructing (9). As a first step, we will construct a chain-level homomorphism

$$\rho_{\widehat{S}} \colon \mathrm{C}_{2g-2}(\mathcal{T}_{2g}(\mathbb{F}_p);\mathbb{Q}) \longrightarrow \mathrm{C}_{2g-k-1}(X;\mathbb{Q}).$$

It is enough to define $\rho_{\widehat{S}}(\eta)$ for a $(2g-2)$-dimensional simplex η of $\mathcal{T}_{2g}(\mathbb{F}_p)$, which we do as follows.

- If the flag associated to η contains $T_1, \ldots, T_{k-1}$, then define $\rho_{\widehat{S}}(\eta)$ to be the $(2g-k-1)$-dimensional simplex of X obtained by deleting $T_1, \ldots, T_{k-1}$ from μ.
- Otherwise, define $\rho_S(\mu) = 0$.

Next, we prove that $\rho_{\widehat{S}}$ induces a map on homology. We remark that this is not obvious – though one could define $\rho_{\widehat{S}}$ on lower-dimensional simplices of $\mathcal{T}_{2g}(\mathbb{F}_p)$ in an analogous manner, the result is not a map of chain complexes. Since $\mathcal{T}_{2g}(\mathbb{F}_p)$ and X are $(2g-2)$-dimensional and $(2g-k-1)$-dimensional, respectively, we have $\mathrm{H}_{2g-2}(\mathcal{T}_{2g}(\mathbb{F}_p);\mathbb{Q}) = \mathrm{Z}_{2g-2}(\mathcal{T}_{2g}(\mathbb{F}_p);\mathbb{Q})$ and $\mathrm{H}_{2g-k-1}(X;\mathbb{Q}) = \mathrm{Z}_{2g-k-1}(X;\mathbb{Q})$. Consider $x \in \mathrm{Z}_{2g-2}(\mathcal{T}_{2g}(\mathbb{F}_p);\mathbb{Q})$. We must show that $\rho_{\widehat{S}}(x) \in \mathrm{Z}_{2g-k-1}(X;\mathbb{Q})$. Letting ζ be a $(2g-k-2)$-simplex of X and letting $\partial_\zeta \colon \mathrm{C}_{2g-k-1}(X;\mathbb{Q}) \to \mathbb{Q}$ be the homomorphism that outputs the coefficient of ζ under the simplicial boundary map, this is equivalent to showing that $\partial_\zeta(\rho_{\widehat{S}}(x)) = 0$. Let $\widetilde{\zeta}$ be the $(2g-3)$-simplex of $\mathcal{T}_{2g}(\mathbb{F}_p)$ obtained by inserting $T_1, \ldots, T_{k-1}$ into the flag associated to ζ and let $\partial_{\widetilde{\zeta}} \colon \mathrm{C}_{2g-2}(\mathcal{T}_{2g}(\mathbb{F}_p);\mathbb{Q}) \to \mathbb{Q}$ be the homomorphism that outputs the coefficient of $\widetilde{\zeta}$ under the simplicial boundary map. It is not necessarily true that $\partial_\zeta \circ \rho_{\widehat{S}} = \partial_{\widetilde{\zeta}}$; however, this is true up to a sign that depends on ζ (the T_i terms that are inserted change the signs in the boundary map). Since $x \in \mathrm{Z}_{2g-2}(\mathcal{T}_{2g}(\mathbb{F}_p);\mathbb{Q})$, we have $\partial_{\widetilde{\eta}}(x) = 0$, which implies that $\partial_\zeta(\rho_{\widehat{S}}(x)) = 0$, as desired.

This completes the construction of $\pi_{\widehat{S}}$. The two indicated properties are immediate from the construction. □

Projection maps, unordered. Let S be a symplectic splitting of $\mathbb{F}_p^{2g}$. We construct a projection map $\pi_S\colon \mathrm{St}_{2g}(\mathbb{F}_p) \to \mathrm{St}(S)$ by averaging the maps constructed in Lemma 4.4 associated to all possible orderings of S. More precisely, let $\widehat{S} = (S_1, \dots, S_k)$ be an ordering of S. We then define

$$\pi_S(x) = \frac{1}{k!} \sum_{\sigma} \pi_{\sigma(\widehat{S})}(x) \qquad (x \in \mathrm{St}_{2g}(\mathbb{F}_p)),$$

where σ ranges over all permutations of $\{1, \dots, k\}$ and $\sigma(\widehat{S}) = (S_{\sigma(1)}, \dots, S_{\sigma(k)})$. This does not depend on the choice of ordering $\widehat{S}$. The one property of π_S that we will need is as follows.

Lemma 4.5. *Let $g \geqslant 1$ and $p \geqslant 2$ be such that p is prime. Let $S = \{S_1, \dots, S_k\}$ be a symplectic splitting of $\mathbb{F}_p^{2g}$ and let $\mathfrak{A}_{B_i}$ be an apartment class in $\mathrm{St}(S_i)$ for $1 \leqslant i \leqslant k$. Then*

$$\pi_S(\mathfrak{A}_{B_1 \cdots B_k}) = \mathit{Sym}(\mathfrak{A}_{B_1}, \dots, \mathfrak{A}_{B_k}).$$

Proof. Immediate from the formulas in Lemma 4.4. □

Corollary 4.6. *Let $g \geqslant 1$ and $p \geqslant 2$ be such that p is prime. Let S be a symplectic splitting of $\mathbb{F}_p^{2g}$. Then ι_S is injective and π_S is surjective.*

Proof. The formulas in Lemmas 4.3 and 4.5 imply that $\pi_S \circ \iota_S = \mathrm{id}$. □

Proof of main theorem. We finally come to the proof of our main theorem.

Proof of Theorem 4.1. We begin by stating a more precise version of Theorem 4.1. Given $S \in \mathfrak{S}_g$, let let $\iota_S^{\mathrm{ns}}\colon \mathrm{St}^{\mathrm{ns}}(S) \to \mathrm{St}_{2g}(\mathbb{F}_p)$ be the restriction to $\mathrm{St}^{\mathrm{ns}}(S)$ of the map $\iota_S\colon \mathrm{St}(S) \to \mathrm{St}_{2g}(\mathbb{F}_p)$ given by Lemma 4.3. Here we are regarding $\mathrm{St}^{\mathrm{ns}}(S)$ as a subspace of $\mathrm{St}(S)$ as in (6). Also, let $\pi_S^{\mathrm{ns}}\colon \mathrm{St}_{2g}(\mathbb{F}_p) \to \mathrm{St}^{\mathrm{ns}}(S)$ be the composition of the map $\pi_S\colon \mathrm{St}_{2g}(\mathbb{F}_p) \to \mathrm{St}(S)$ defined above with the quotient map $\mathrm{St}(S) \to \mathrm{St}^{\mathrm{ns}}(S)$. Finally, define

$$\iota^{\mathrm{ns}}\colon \bigoplus_{S \in \mathfrak{S}_g} \mathrm{St}^{\mathrm{ns}}(S) \longrightarrow \mathrm{St}_{2g}(\mathbb{F}_p)$$

to be the direct sum of all the ι_S^{ns} and define

$$\pi^{\mathrm{ns}}\colon \mathrm{St}_{2g}(\mathbb{F}_p) \longrightarrow \bigoplus_{S \in \mathfrak{S}_g} \mathrm{St}^{\mathrm{ns}}(S)$$

to be the direct product of all the π_S^{ns}. We will prove that ι^{ns} and π^{ns} are isomorphisms satisfying $\iota^{\mathrm{ns}} \circ \pi^{\mathrm{ns}} = \mathrm{id}$ and $\pi^{\mathrm{ns}} \circ \iota^{\mathrm{ns}} = \mathrm{id}$. To do this, it is enough to prove that ι^{ns} is surjective and that $\pi^{\mathrm{ns}} \circ \iota^{\mathrm{ns}} = \mathrm{id}$, which are the content of Steps 1 and 5 below.

Before we launch into the proof, we introduce some terminology. Say that $S' \in \mathfrak{S}_g$ is a *refinement* of $S \in \mathfrak{S}_g$ if each subspace occurring in S is the direct sum of a subset of the subspaces occurring in S'. This gives a partial order on $\mathfrak{S}_g$ where $S' \preceq S$ when S' is a refinement of S.

Step 1. *We prove that ι^{ns} is surjective.*

Proof of Step 1. Since permuting the vectors making up an apartment class only changes the apartment class by a sign (see the second relation in Theorem 3.6), for $S, S' \in \mathfrak{S}_g$ such that $S' \preceq S$ we have

$$\iota_{S'}(\mathrm{St}(S')) \subset \iota_S(\mathrm{St}(S)).$$

In fact, even more is true: examining the definitions, we see that for $S \in \mathfrak{S}_g$ we have

$$\iota_S(\mathrm{St}^{\mathrm{sep}}(S)) = \sum_{S' \prec S} \iota_{S'}(\mathrm{St}(S')). \tag{10}$$

We now prove a generalization of the fact that ι^{ns} is surjective. Namely, we will prove that for all $S \in \mathfrak{S}_g$, we have

$$\iota_S(\mathrm{St}(S)) = \sum_{S' \preceq S} \iota_{S'}(\mathrm{St}^{\mathrm{ns}}(S')). \tag{11}$$

The assertion that ι^{ns} is surjective is equivalent to this for S equal to the 1-element symplectic splitting $(\mathbb{F}_p^{2g})$. Consider $S \in \mathfrak{S}_g$. We can assume by induction that (11) holds for all smaller S (the base case, when S is minimal under $\preceq$, is vacuous). We have $\mathrm{St}(S) = \mathrm{St}^{\mathrm{ns}}(S) \oplus \mathrm{St}^{\mathrm{sep}}(S)$, which implies that

$$\begin{aligned}
\iota_S(\mathrm{St}(S)) &= \iota_S(\mathrm{St}^{\mathrm{ns}}(S)) + \iota_S(\mathrm{St}^{\mathrm{sep}}(S)) \\
&= \iota_S(\mathrm{St}^{\mathrm{ns}}(S)) + \sum_{S' \prec S} \iota_{S'}(\mathrm{St}(S')) \\
&= \iota_S(\mathrm{St}^{\mathrm{ns}}(S)) + \sum_{S' \prec S} \iota_{S'}(\mathrm{St}^{\mathrm{ns}}(S')).
\end{aligned}$$

The second equality follows from (10) and the third equality follows from our inductive hypothesis. The assertion (11) follows. □

Our next goal is to study the map $\pi_{S'} \circ \iota_S \colon \mathrm{St}(S) \to \mathrm{St}(S')$ for $S, S' \in \mathfrak{S}_g$. We divide this into several cases.

Step 2. *For $S \in \mathfrak{S}_g$, the map $\pi_S \circ \iota_S \colon \mathrm{St}(S) \to \mathrm{St}(S)$ is the identity. Consequently, $\pi_S^{ns} \circ \iota_S^{ns} = id$.*

Proof of Step 2. Immediate from the formulas in Lemmas 4.3 and 4.5. □

Step 3. *Consider $S, S' \in \mathfrak{S}_g$. Assume that S' is not a refinement of S. Then the image of $\pi_{S'} \circ \iota_S \colon \mathrm{St}(S) \to \mathrm{St}(S')$ is contained in $\mathrm{St}^{sep}(S')$. Consequently, $\pi_{S'}^{ns} \circ \iota_S^{ns} = 0$.*

Proof of Step 3. Let $\widehat{S}'$ be any ordering of S'. It is enough to prove that the image of $\pi_{\widehat{S}'} \circ \iota_S$ is contained in $\mathrm{St}^{\mathrm{sep}}(S')$. Write $S = \{S_1, \ldots, S_k\}$ and $\widehat{S}' = (S'_1, \ldots, S'_\ell)$. By the formulas in Lemmas 4.3 and 4.4, it is enough to prove the following.

- Let B be an ordered basis for $\mathbb{F}_p^{2g}$ such that each vector in B lies in some S_i. Assume that B is weakly compatible with $\widehat{S}'$ with factorization $B = B_1 \cdots B_\ell$ and projection $(\overline{B}_1, \ldots, \overline{B}_\ell)$. Then
$$\mathrm{Sym}(\mathfrak{A}_{\overline{B}_1}, \ldots, \mathfrak{A}_{\overline{B}_\ell}) \in \mathrm{St}^{\mathrm{sep}}(S').$$

Since permuting the vectors in an apartment class only changes the apartment class by a sign (see Theorem 3.6), we can permute the vectors in each B_i freely. Doing this, we can assume that each B_i is of the form $B_i = B_{i1} \cdots B_{ik}$, where B_{ij} is a sequence of vectors (possibly the empty sequence) lying in S_j. For $1 \leqslant i \leqslant \ell$ and $1 \leqslant j \leqslant k$, let $\overline{B}_{ij}$ be the image of B_{ij} in $\overline{B}_i$. We thus have $\overline{B}_i = \overline{B}_{i1} \cdots \overline{B}_{ik}$.

For $1 \leqslant i \leqslant \ell$ and $1 \leqslant j \leqslant k$, let T_{ij} be the span of the vectors in B_{ij} and let $\overline{T}_{ij}$ be the span of the vectors in $\overline{B}_{ij}$. By construction, we have

$$T_{ij} \subset S'_1 \oplus \cdots \oplus S'_i \quad \text{and} \quad T_{ij} \subset S_j.$$

Also, we have $\overline{T}_{ij} \subset S'_i$; however, in projecting T_{ij} into S'_i we might have destroyed the fact that it lies in S_j. However, we claim that this is not the case and $\overline{T}_{ij} \subset S_j$. To see this, observe that by definition, $\overline{T}_{ij}$ is the image of T_{ij} under the projection map

$$\rho_{ij} \colon S'_1 \oplus \cdots \oplus S'_{i-1} \oplus S'_i \longrightarrow S'_i.$$

The domain of ρ_{ij} is equal to

$$\left(\bigoplus_{\substack{1 \leqslant i' \leqslant i-1 \\ 1 \leqslant j' \leqslant k}} T_{i'j'} \right) \oplus S'_i.$$

We have $T_{ij} \subset S_j$ and $T_{i'j'} \subset S_{j'}$. If $j \neq j'$, then the subspaces S_j and $S_{j'}$ are orthogonal. It follows that the projection of $x \in T_{ij}$ to $T_{i'j'}$ is 0 unless $j = j'$. We deduce that for $x \in T_{ij} \subset S_j$, we have $\rho_{ij}(x) = x - y$, where

$$y \in \bigoplus_{i'=1}^{i-1} T_{i'j} \subset S_j.$$

This implies that $\rho_{ij}(x) \in S_j$, so $\overline{T}_{ij} = \rho_{ij}(T_{ij}) \subset S_j$, as claimed.

Define

$$\overline{T} = \{\overline{T}_{ij} \mid 1 \leqslant i \leqslant \ell,\, 1 \leqslant j \leqslant k,\, \overline{T}_{ij} \neq 0\}.$$

Since $\overline{T}_{ij} \subset S'_i \cap S_j$ and both S and S' are symplectic splittings, it follows that the elements of $\overline{T}$ are pairwise orthogonal to each other. They also span $\mathbb{F}_p^{2g}$, so we deduce that $\overline{T}$ is a symplectic splitting of $\mathbb{F}_p^{2g}$. We have $\overline{T} \preceq S$ and $\overline{T} \preceq S'$; since S' is not a refinement of S, we must have $\overline{T} \prec S'$. This implies that there exists some $1 \leqslant i \leqslant \ell$ such that

$$\overline{T}_i \coloneqq \{\overline{T}_{ij} \mid 1 \leqslant j \leqslant k,\, \overline{T}_{ij} \neq 0\} \neq \{S'_i\}.$$

Since $\overline{T}_{ij}$ is the span of the vectors in $\overline{B}_{ij}$ and $\overline{B}_i = \overline{B}_{i1} \cdots \overline{B}_{ik}$, we deduce that $\mathfrak{A}_{\overline{B}_i} \in \mathrm{St}^{\mathrm{sep}}(S'_i)$ and thus that

$$\mathrm{Sym}(\mathfrak{A}_{\overline{B}_1}, \ldots, \mathfrak{A}_{\overline{B}_\ell}) \in \mathrm{St}^{\mathrm{sep}}(S'),$$

as desired. □

The next step requires a definition. For $S \in \mathfrak{S}_g$, define

$$\iota^{\mathrm{ns}}_{\prec S} \colon \bigoplus_{S' \prec S} \mathrm{St}^{\mathrm{ns}}(S') \longrightarrow \mathrm{St}_{2g}(\mathbb{F}_p)$$

to be the direct sum of all the $\iota^{\mathrm{ns}}_{S'}$ for $S' \prec S$ and define

$$\pi^{\mathrm{ns}}_{\prec S} \colon \mathrm{St}_{2g}(\mathbb{F}_p) \longrightarrow \bigoplus_{S' \prec S} \mathrm{St}^{\mathrm{ns}}(S')$$

to be the direct product of all the $\pi^{\mathrm{ns}}_{S'}$ for $S' \prec S$.

Step 4. *For $S \in \mathfrak{S}_g$, the kernel of*

$$\pi^{ns}_{\prec S} \circ \iota_S \colon \mathrm{St}(S) \longrightarrow \bigoplus_{S' \prec S} \mathrm{St}^{ns}(S')$$

is $\mathrm{St}^{sep}(S)$. Consequently, if $S' \in \mathfrak{S}_g$ is a refinement of $S \in \mathfrak{S}_g$, then $\pi^{ns}_{S'} \circ \iota^{ns}_{S} = 0$.

Proof of Step 4. We can assume by induction that the step is true for all symplectic splittings that are smaller than S in the partial order on symplectic splittings (a vacuous assumption if S is minimal). Corollary 4.6 says that ι_S is injective. Using this, we will identify $\mathrm{St}(S)$ with its image in $\mathrm{St}_{2g}(\mathbb{F}_p)$. The proof of Step 1 shows that the image of $\iota^{\mathrm{ns}}_{\prec S}$ is $\mathrm{St}^{\mathrm{sep}}(S) \subset \mathrm{St}(S)$. What is more, we can combine Steps 2 and 3 with our inductive hypothesis to see that for $S', S'' \prec S$ we have

$$\pi^{\mathrm{ns}}_{S''} \circ \iota^{\mathrm{ns}}_{S'} = \begin{cases} \mathrm{id} & \text{if } S' = S'', \\ 0 & \text{otherwise.} \end{cases}$$

This implies that $\pi^{\mathrm{ns}}_{\prec S} \circ \iota^{\mathrm{ns}}_{\prec S} = \mathrm{id}$, and thus that $\iota^{\mathrm{ns}}_{\prec S}$ is injective and $\pi^{\mathrm{ns}}_{\prec S}$ is surjective. Since both $\pi^{\mathrm{ns}}_{\prec S}$ and $\iota^{\mathrm{ns}}_{\prec S}$ are $\mathrm{Sp}(S)$-equivariant, we can identify the domain of $\iota^{\mathrm{ns}}_{\prec S}$ and the codomain of $\pi^{\mathrm{ns}}_{\prec S}$ with $\mathrm{St}^{\mathrm{sep}}(S)$. Using this identification, the restriction of $\pi^{\mathrm{ns}}_{\prec S}$ to $\mathrm{St}(S)$ is precisely the projection to $\mathrm{St}^{\mathrm{sep}}(S)$ whose kernel is $\mathrm{St}^{\mathrm{ns}}(S)$ (see the decomposition (6)). The step follows. □

All the pieces are now in place to prove that $\pi^{\mathrm{ns}} \circ \iota^{\mathrm{ns}} = \mathrm{id}$.

Step 5. *We prove that $\pi^{ns} \circ \iota^{ns} = id$.*

Proof of Step 5. Steps 2, 3, and 4 say that for $S, S' \in \mathfrak{S}_g$ we have

$$\pi^{\mathrm{ns}}_{S'} \circ \iota^{\mathrm{ns}}_{S} = \begin{cases} \mathrm{id} & \text{if } S = S', \\ 0 & \text{otherwise.} \end{cases}$$

The step follows. □

This completes the proof of Theorem 4.1. □

5 Bounds

In this section, we prove Proposition 3.9. We first recall its statement. Fix a prime p. For $g \geqslant 1$, define

$$\theta_g = \frac{\dim_{\mathbb{Q}} \mathrm{St}^{\mathrm{ns}}_{2g}(\mathbb{F}_p)}{|\mathrm{Sp}_{2g}(\mathbb{F}_p)|} \quad \text{and} \quad \lambda_g = \frac{1}{g(p^{2g}-1)}.$$

Proposition 3.9 asserts that $\theta_g = \lambda_g$ for all $g \geqslant 1$.

Before we prove this, we must prove two lemmas that assert that the θ_g and λ_g satisfy similar recurrence relations. Given $n \geqslant 1$, a *partition* of n is an expression

$$n = n_1 + n_2 + \cdots + n_m \tag{12}$$

with

$$n_1 \geqslant n_2 \geqslant \cdots \geqslant n_m \geqslant 1.$$

If the numbers that appear in (12) are $a_1 > a_2 > \cdots > a_\ell$ and a_i appears r_i times, then we will denote the partition by $(a_1^{r_1}, \ldots, a_\ell^{r_\ell})$ and say as shorthand that

$$(a_1^{r_1}, \ldots, a_\ell^{r_\ell}) \vdash n.$$

Our first lemma is then as follows.

Lemma 5.1. *For $g \geqslant 1$ and $p \geqslant 2$ such that p is prime, we have*

$$\frac{p^{\binom{2g}{2}}}{|\operatorname{Sp}_{2g}(\mathbb{F}_p)|} = \sum_{(a_1^{r_1}, \ldots, a_\ell^{r_\ell}) \vdash g} \frac{(\theta_{a_1})^{r_1}(\theta_{a_2})^{r_2} \cdots (\theta_{a_\ell})^{r_\ell}}{(r_1)!(r_2)! \cdots (r_\ell)!}.$$

Proof. We begin with some combinatorial facts about symplectic splittings. Given an ordered symplectic splitting $\widehat{S} = (S_1, \ldots, S_k)$ of $\mathbb{F}_p^{2g}$ and a partition $(a_1^{r_1}, \ldots, a_\ell^{r_\ell}) \vdash g$, we say that the *type* of $\widehat{S}$ is $(a_1^{r_1}, \ldots, a_\ell^{r_\ell})$ if the sequence of numbers

$$(\dim_{\mathbb{F}_p}(S_1), \dim_{\mathbb{F}_p}(S_2), \ldots, \dim_{\mathbb{F}_p}(S_k))$$

begins with r_1 entries of a_1, then has r_2 entries of a_2, etc. We claim that there are

$$\frac{|\operatorname{Sp}_{2g}(\mathbb{F}_p)|}{\prod_{i=1}^{\ell} |\operatorname{Sp}_{2a_i}(\mathbb{F}_p)|^{r_i}} \tag{13}$$

ordered symplectic splittings of type $(a_1^{r_1}, \ldots, a_\ell^{r_\ell})$. Indeed, the group $\operatorname{Sp}_{2g}(\mathbb{F}_p)$ acts transitively on the ordered symplectic splittings of type $(a_1^{r_1}, \ldots, a_\ell^{r_\ell})$ and $\prod_{i=1}^{\ell} \operatorname{Sp}_{2a_i}(\mathbb{F}_p)^{r_i}$ is the stabilizer of one such splitting.

Next, we say that an (unordered) symplectic splitting has type $(a_1^{r_1}, \ldots, a_\ell^{r_\ell})$ if some ordering of it has that type. We claim that there are

$$\frac{|\operatorname{Sp}_{2g}(\mathbb{F}_p)|}{\prod_{i=1}^{\ell} (r_i)! |\operatorname{Sp}_{2a_i}(\mathbb{F}_p)|^{r_i}} \tag{14}$$

symplectic splittings of type $(a_1^{r_1}, \ldots, a_\ell^{r_\ell})$. Indeed, there are $(r_1)!(r_2)! \cdots (r_\ell)!$ orderings of that type for each unordered splitting, so the equation follows from (13) above.

We now turn to the desired identity. Theorem 4.1 implies that

$$\dim_{\mathbb{Q}} \operatorname{St}_{2g}(\mathbb{F}_p) = \sum_{S \in \mathfrak{S}_g} \dim_{\mathbb{Q}} \operatorname{St}^{\mathrm{ns}}(S). \tag{15}$$

The Solomon–Tits theorem [Brow, Theorem IV.5.2] implies that $\dim_{\mathbb{Q}} \operatorname{St}_{2g}(\mathbb{F}_p) = p^{\binom{2g}{2}}$. Also, by definition for $S \in \mathfrak{S}_g$ with $S = \{S_1, \ldots, S_k\}$ we have $\dim_{\mathbb{Q}} \operatorname{St}^{\mathrm{ns}}(S) = \prod_{i=1}^{k} \operatorname{St}^{\mathrm{ns}}(S_i)$. Plugging these into (15) and grouping the elements of $\mathfrak{S}_g$ together by type, we can use (14) to deduce that

$$\begin{aligned} p^{\binom{2g}{2}} &= \sum_{(a_1^{r_1}, \ldots, a_\ell^{r_\ell}) \vdash g} \frac{|\operatorname{Sp}_{2g}(\mathbb{F}_p)| (\dim_{\mathbb{Q}} \operatorname{St}^{\mathrm{ns}}_{2a_1})^{r_1} (\dim_{\mathbb{Q}} \operatorname{St}^{\mathrm{ns}}_{2a_2})^{r_2} \cdots (\dim_{\mathbb{Q}} \operatorname{St}^{\mathrm{ns}}_{2a_\ell})^{r_\ell}}{\prod_{i=1}^{\ell} (r_i)! |\operatorname{Sp}_{2a_i}(\mathbb{F}_p)|^{r_i}} \\ &= \sum_{(a_1^{r_1}, \ldots, a_\ell^{r_\ell}) \vdash g} \frac{|\operatorname{Sp}_{2g}(\mathbb{F}_p)| (\theta_{a_1})^{r_1} (\theta_{a_2})^{r_2} \cdots (\theta_{a_\ell})^{r_\ell}}{(r_1)!(r_2)! \cdots (r_\ell)!}. \end{aligned}$$

Dividing both sides by $|\operatorname{Sp}_{2g}(\mathbb{F}_p)|$ gives the desired result. ∎

Our second lemma is a similar identity for $\lambda_g = \frac{1}{g(p^{2g}-1)}$. Its proof was suggested to the authors by the MathOverflow user "Lucia"; see [Lu].

Lemma 5.2. *For $g \geqslant 1$ and $p \geqslant 2$ such that p is prime, we have*

$$\frac{p^{\binom{2g}{2}}}{|\operatorname{Sp}_{2g}(\mathbb{F}_p)|} = \sum_{(a_1^{r_1},\dots,a_\ell^{r_\ell})\vdash g} \frac{(\lambda_{a_1})^{r_1}(\lambda_{a_2})^{r_2}\cdots(\lambda_{a_\ell})^{r_\ell}}{(r_1)!(r_2)!\cdots(r_\ell)!}.$$

Proof. Using the standard formula for $|\operatorname{Sp}_{2g}(\mathbb{F}_p)|$ (see, e.g., [G, p. 27]), we see that the left hand side of the purported identity equals

$$\begin{aligned}\frac{p^{g(2g-1)}}{(p^{2g}-1)p^{2g-1}(p^{2g-2}-1)p^{2g-3}\cdots(p^2-1)p} &= \frac{p^{g(2g-1)}}{p^{g^2}(p^{2g}-1)(p^{2g-2}-1)\cdots(p^2-1)} \\ &= \frac{p^{g(g-1)}}{(p^{2g}-1)(p^{2g-2}-1)\cdots(p^2-1)}.\end{aligned}$$

We now focus on the right hand side. Form the generating function

$$\begin{aligned}\Lambda(x) &= \sum_{g=1}^{\infty}\left(\sum_{(a_1^{r_1},\dots,a_\ell^{r_\ell})\vdash g}\frac{(\lambda_{a_1})^{r_1}(\lambda_{a_2})^{r_2}\cdots(\lambda_{a_\ell})^{r_\ell}}{(r_1)!(r_2)!\cdots(r_\ell)!}\right)x^g \\ &= \sum_{g=1}^{\infty}\left(\sum_{(a_1^{r_1},\dots,a_\ell^{r_\ell})\vdash g}\frac{(p^{2a_1}-1)^{-r_1}(p^{2a_2}-1)^{-r_2}\cdots(p^{2a_\ell}-1)^{-r_\ell}}{(a_1)^{r_1}(r_1)!(a_2)^{r_2}(r_2)!\cdots(a_\ell)^{r_\ell}(r_\ell)!}\right)x^g.\end{aligned}$$

The first observation is that the Exponential Formula [St, Corollary 5.1.9] implies that

$$\Lambda(x) = \exp\left(\sum_{g=1}^{\infty}\left(\frac{(p^{2g}-1)^{-1}}{g}\right)x^g\right) = \exp\left(\sum_{g=1}^{\infty}\frac{x^g}{g(p^{2g}-1)}\right).$$

See [Y] and the references therein for more details. We now compute that

$$\begin{aligned}\exp\left(\sum_{g=1}^{\infty}\frac{x^g}{g(p^{2g}-1)}\right) &= \exp\left(-\sum_{g=1}^{\infty}\left(\frac{x^g}{g}\sum_{h=0}^{\infty}p^{2gh}\right)\right) = \exp\left(-\sum_{h=0}^{\infty}\sum_{g=1}^{\infty}\frac{(xp^{2h})^g}{g}\right) \\ &= \exp\left(\sum_{h=0}^{\infty}\log(1-p^{2h}x)\right) = \prod_{h=0}^{\infty}(1-p^{2h}x).\end{aligned}$$

A theorem of Euler (see [A, Corollary 2.2]) asserts that

$$\begin{aligned}\prod_{h=0}^{\infty}(1-p^{2h}x) &= 1+\sum_{g=1}^{\infty}\left(\frac{(-1)^g p^{g(g-1)}}{(1-p^{2g})(1-p^{2g-2})\cdots(1-p^2)}\right)x^g \\ &= 1+\sum_{g=1}^{\infty}\left(\frac{p^{g(g-1)}}{(p^{2g}-1)(p^{2g-2}-1)\cdots(p^2-1)}\right)x^g.\end{aligned}$$

The coefficient of x^g in this equals the above expression for the left hand side of our purported identity. The lemma follows. ∎

Proof of Proposition 3.9. We will prove that $\theta_g = \lambda_g$ for $g \geqslant 1$ by induction on g. The base case $g = 1$ asserts that

$$\frac{\dim_{\mathbb{Q}} \mathrm{St}_2^{\mathrm{ns}}(\mathbb{F}_p)}{|\,\mathrm{Sp}_2(\mathbb{F}_p)|} = \frac{1}{(p^2-1)}.$$

Since there are no nontrivial symplectic splittings of $\mathbb{F}_p^2$, we have $\mathrm{St}_2^{\mathrm{ns}}(\mathbb{F}_p) = \mathrm{St}_2(\mathbb{F}_p)$. The Solomon–Tits Theorem [Brow, Theorem IV.5.2] says that $\mathrm{St}_2(\mathbb{F}_p)$ is p-dimensional. What is more, $|\,\mathrm{Sp}_2(\mathbb{F}_p)| = (p^2-1)p$ (see, e.g., [G, p. 27]). The above identity follows.

Now assume that $g > 1$ and that $\theta_{g'} = \lambda_{g'}$ for all $1 \leqslant g' < g$. Applying Lemmas 5.1 and 5.2 together with our inductive hypothesis, we see that

$$\begin{aligned}\theta_g &= \frac{p^{\binom{2g}{2}}}{|\,\mathrm{Sp}_{2g}(\mathbb{F}_p)|} - \sum_{\substack{(a_1^{r_1},\ldots,a_\ell^{r_\ell})\vdash g \\ \ell\geqslant 2}} \frac{(\theta_{a_1})^{r_1}(\theta_{a_2})^{r_2}\cdots(\theta_{a_\ell})^{r_\ell}}{(r_1)!(r_2)!\cdots(r_\ell)!} \\ &= \frac{p^{\binom{2g}{2}}}{|\,\mathrm{Sp}_{2g}(\mathbb{F}_p)|} - \sum_{\substack{(a_1^{r_1},\ldots,a_\ell^{r_\ell})\vdash g \\ \ell\geqslant 2}} \frac{(\lambda_{a_1})^{r_1}(\lambda_{a_2})^{r_2}\cdots(\lambda_{a_\ell})^{r_\ell}}{(r_1)!(r_2)!\cdots(r_\ell)!} = \lambda_g.\end{aligned}$$

□

References

[A] G. E. Andrews, *The theory of partitions*, Addison-Wesley Publishing Co., Reading, MA, 1976.

[Broa] N. Broaddus, Homology of the curve complex and the Steinberg module of the mapping class group, Duke Math. J. 161 (2012), no. 10, 1943–1969.

[Brow] K. S. Brown, *Buildings*, Springer, New York, 1989.

[Cha] C. Chaudhuri, The cohomological excess of certain moduli spaces of curves of genus g, Int. Math. Res. Not. IMRN 2015, no. 4, 1056–1074.

[ChuFarPu] T. Church, B. Farb and A. Putman, A stability conjecture for the unstable cohomology of $\mathrm{SL}_n(\mathbb{Z})$, mapping class groups, and $\mathrm{Aut}(F_n)$, in *Algebraic topology: applications and new directions*, 55–70, Contemp. Math., 620, Amer. Math. Soc., Providence, RI.

[FabLo] C. Faber and E. Looijenga, Remarks on moduli of curves, in *Moduli of curves and abelian varieties*, 23–45, Aspects Math., E33, Vieweg, Braunschweig.

[FarMar] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, 49, Princeton Univ. Press, Princeton, NJ, 2012.

[FoPa] C. Fontanari and S. Pascolutti, An affine open covering of M_g for $g \leqslant 5$, Geom. Dedicata 158 (2012), 61–68.

[G] L. C. Grove, *Classical groups and geometric algebra*, Graduate Studies in Mathematics, 39, Amer. Math. Soc., Providence, RI, 2002.

[Hai] R. M. Hain, Torelli groups and geometry of moduli spaces of curves, in *Current topics in complex algebraic geometry (Berkeley, CA, 1992/93)*, 97–143, Math. Sci. Res. Inst. Publ., 28, Cambridge Univ. Press, Cambridge.

[Hare1] J. L. Harer, Stability of the homology of the mapping class groups of orientable surfaces, Ann. of Math. (2) 121 (1985), no. 2, 215–249.

[Hare2] J. L. Harer, The virtual cohomological dimension of the mapping class group of an orientable surface, Invent. Math. 84 (1986), no. 1, 157–176.

[HareZ] J. Harer and D. Zagier, The Euler characteristic of the moduli space of curves, Invent. Math. 85 (1986), no. 3, 457–485.

[Hart] R. Hartshorne, Cohomological dimension of algebraic varieties, Ann. of Math. (2) 88 (1968), 403–450.

[Hat] A. Hatcher, On triangulations of surfaces, Topology Appl. 40 (1991), no. 2, 189–194.

[HatV] A. Hatcher and K. Vogtmann, Tethers and homology stability for surfaces, 2015 preprint.

[KLeiSc] R. P. Kent, IV, C. J. Leininger and S. Schleimer, Trees and mapping class groups, J. Reine Angew. Math. 637 (2009), 1–21.

[LeeSz] R. Lee and R. H. Szczarba, On the homology and cohomology of congruence subgroups, Invent. Math. 33 (1976), no. 1, 15–53.

[Lu] Lucia (http://mathoverflow.net/users/38624/lucia), Identity involving a sum over all partitions of n, URL (version: 2016-09-27): http://mathoverflow.net/q/250862

[MadW] I. Madsen and M. Weiss, The stable moduli space of Riemann surfaces: Mumford's conjecture, Ann. of Math. (2) 165 (2007), no. 3, 843–941.

[Mon] G. Mondell, On the cohomological dimension of the moduli space of Riemann surfaces, to appear in Duke Math. J.

[MorSaSu] S. Morita, T. Sakasai and M. Suzuki, Abelianizations of derivation Lie algebras of the free associative algebra and the free Lie algebra, Duke Math. J. 162 (2013), no. 5, 965–1002.

[Pu] A. Putman, The second rational homology group of the moduli space of curves with level structures, Adv. Math. 229 (2012), no. 2, 1205–1234.

[Se] J.-P. Serre, Sur la cohomologie des variétés algébriques, J. Math. Pures Appl. (9) 36 (1957), 1–16.

[So] L. Solomon, The Steinberg character of a finite group with BN-pair, in *Theory of Finite Groups (Symposium, Harvard Univ., Cambridge, Mass., 1968)*, 213–221, Benjamin, New York.

[St] R. P. Stanley, *Enumerative combinatorics. Vol. 2*, Cambridge Studies in Advanced Mathematics, 62, Cambridge Univ. Press, Cambridge, 1999.

[Y] Qiaochu Yuan (http://mathoverflow.net/users/290/qiaochu-yuan), A remarkable sum over partitions, URL (version: 2015-11-21): http://mathoverflow.net/q/224159

Neil Fullarton
Department of Mathematics
Rice University, MS 136
6100 Main St.
Houston, TX 77005
E-mail: neil.fullarton@rice.edu

Andrew Putman
Department of Mathematics
University of Notre Dame
279 Hurley Hall
Notre Dame, IN 46556
E-mail: andyp@nd.edu